

Semantic and Associative Priming in a Distributed Attractor Network

David C. Plaut

Department of Psychology
Carnegie Mellon University, and
Center for the Neural Basis of Cognition
Pittsburgh, PA 15213–3890
plaut@cmu.edu

Abstract

A distributed attractor network is trained on an abstract version of the task of deriving the meanings of written words. When processing a word, the network starts from the final activity pattern of the previous word. Two words are semantically related if they overlap in their semantic features, whereas they are associatively related if one word follows the other frequently during training. After training, the network exhibits two empirical effects that have posed problems for distributed network theories: much stronger associative priming than semantic priming, and significant associative priming across an intervening unrelated item. It also reproduces the empirical findings of greater priming for low-frequency targets, degraded targets, and high-dominance category exemplars.

Introduction

In a variety of lexical tasks, including naming and lexical decision, subjects are faster and more accurate to process a word, such as BUTTER, when it is preceded by a semantically related word, such as BREAD (see Neely, 1991, for a review). Such findings of semantic priming are taken by many theorists as reflecting fundamental properties of the organization of knowledge in the human cognitive system.

Broadly speaking, two classes of theories of semantic memory have been put forth to account for semantic priming in lexical tasks. Spreading-activation theories (e.g., Anderson, 1983; Collins & Loftus, 1975; McNamara, 1992a, 1992b, 1994) propose that semantic memory consists of a network of interconnected nodes, each representing a particular concept. Processing a word involves activating the concept node in semantic memory corresponding to its meaning. This activation is assumed to spread along links to other, related concepts, thereby facilitating the subsequent processing of those concepts. By contrast, compound-cue theories (e.g., Doshier & Rosedale, 1989; McKoon & Ratcliff, 1992; Ratcliff & McKoon, 1988, 1994) propose that, in processing a word, semantic memory is accessed using a cue consisting of the word conjoined with the context in which it occurs (e.g., the preceding word). Semantically related words co-occur more frequently than do unrelated words, and so the compound cues for related words tend to have greater familiarity than do those for unrelated words. In many general models of memory retrieval (e.g., Gillund & Shiffrin, 1984; Hintzman, 1986; Murdock, 1982), greater familiarity gives rise to faster and more accurate processing, resulting in semantic priming.

Recently, a third type of theory has been proposed to account for semantic priming, based on distributed connectionist networks (Kawamoto, 1988; Masson, 1991, 1995; McRae,

de Sa, & Seidenberg, 1993; Sharkey & Sharkey, 1992). In such networks, each concept is represented, not by a particular unit, but by a particular *pattern* of activity over a large number of processing units. Related concepts are represented by similar (i.e., overlapping) patterns of activity. Each unit can be thought of as encoding a particular semantic feature that participates in many concepts (Smith & Medin, 1981), although these features need not correspond to verbalizable attributes of any concept. In processing a word, units cooperate and compete across weighted connections until the network as a whole settles into a stable pattern of activity that represents the meaning of the word. If the network starts from this pattern in processing a subsequent word, it will be faster to settle for a related than for an unrelated word because many of the units will already be in their correct states.

Distributed network theories bear an interesting relationship both to spreading-activation theories and to compound-cue theories. Although, in some sense, activation “spreads” among units in a distributed network, this spread is not between concepts but between features. For a given pattern of activity, the degree to which any concept is “active” depends on its overlap with the current pattern. After settling into the meaning of a word, related meanings are active *simultaneously* (to the degree that they are similar)—no additional spread of activation is required. Thus, distributed network theories provide a more natural interpretation of the finding that, while the degree of relatedness influences the magnitude of priming, it does not influence its time of onset, which is essentially instantaneous (Lorch, 1982; Ratcliff & McKoon, 1981). In some ways, distributed network theories are more similar to compound-cue theories, in that the processing of a word is sensitive to the context in which it occurs (as reflected in the current state of the network). However, conjunctions need not be represented explicitly, so that processing can depend on properties of contexts independent of their co-occurrences with target words. Thus, for instance, distributed network theories can account naturally for the finding that different neutral contexts (e.g., unrelated words, neutral words like READY, and nonwords) have equivalent effects on between-trial priming, even though their familiarities as compounds with target words are very different (McNamara, 1994).

Unfortunately, two sets of empirical findings appear to pose problems for distributed network theories of semantic priming. The first relates to findings that different types of relations among words influence priming in different ways. In particular, one can distinguish an *associative* relation among words (e.g., as measured by free association norms; Post-

man & Keppel, 1970) from a purely *semantic* relation (i.e., having similar meanings, such as category co-ordinates). Associatively related word pairs are typically also semantically related (e.g., BREAD–BUTTER) but many semantically related word pairs are not associatively related (e.g., BREAD–CAKE). In the few studies that have studied priming among words that are semantically but not associatively related (e.g., Fischler, 1977; Seidenberg, Waters, Sanders, & Langer, 1984), the priming effect is much smaller than that found for associatively related words, particularly in lexical decision, and the effect was completely absent under conditions which prevent expectancies and post-lexical checking (Shelton & Martin, 1992). Furthermore, unlike semantic priming, associative priming is highly asymmetric; for example, BED–PAN produces strong priming whereas PAN–BED produces little if any (particularly in naming; see Neely, 1991). Similar dissociations between semantic and associative effects have been demonstrated using cross-model priming in sentence contexts (Moss & Marslen-Wilson, 1993) and in the degeneration of semantic memory in Alzheimer’s disease (Glosser & Friedman, 1991). These data are problematic for distributed network theories (e.g., Kawamoto, 1988; Masson, 1995) because they typically employ only a single, symmetric manipulation—pattern overlap—to encode word relatedness. There is no opportunity for different types of relations among words—semantic and associative—to behave differently.

The second set of findings that has challenged distributed network theories is that associative priming can span an intervening unrelated item, such as in the word sequence BREAD–DOG–BUTTER, although it is very weak under these conditions (Joordens & Besner, 1992; McNamara, 1994). In a distributed network, if the entire network settles completely to the meaning of the intervening word DOG, then the pattern of activity representing the meaning of BREAD will be completely eliminated, leaving no opportunity for it to facilitate the processing of BUTTER. Masson (1995) considered the possibility that the intervening word might be processed only partially, leaving residual semantic activation from BREAD to influence BUTTER. Using a Hopfield (1982) network, he simulated the small priming effect across unrelated words in a naming task by basing the network’s response on the activity of phonological units which were updated more frequently than semantic units. Unfortunately, the simulations used a very small vocabulary (only three pairs of semantically related items), and no independent justification was provided for why phonological and semantic units should behave differently.

The current paper presents a distributed network model of priming that addresses the challenges posed by associative vs. semantic priming and by priming across an unrelated item. The model differs from previous ones in two main ways. First, whereas semantic relatedness among words is encoded by the degree of overlap of their semantic feature representations, an association from one word to another is encoded directly in the likelihood that the one follows the other during training (see Moss, Hare, Day, & Tyler, 1994, for a similar approach). Second, a more powerful learning procedure is used—continuous back-propagation through time (Pearlmutter, 1989). This procedure has the critical properties that unit states change gradually over time in response to input, and that learning is sensitive to the entire trajectory from the ini-

tial activity pattern to the final activity pattern. The model also replicates a number of basic findings in the priming literature, including greater priming for low-frequency targets, degraded targets, and high-dominance category exemplars.

Simulation

Method

Task. Given that semantic and associative priming have been demonstrated in a wide range of lexical tasks, the current work investigated priming in a general version of the task of understanding words: an abstract version of the task of mapping from written words to their meanings.

The semantic representations of words were generated artificially but with considerable structure. Eight different random patterns were generated over 100 semantic features, in which each unit had a probability of 0.1 of being active. These patterns served as the “prototypes” for eight separate semantic categories. Sixteen category exemplars were generated from each prototype pattern by randomly altering some of its features (Chauvin, 1988). Eight of these were typical or high-dominance exemplars in which relatively few features of the prototype were changed (each feature had a probability of 0.2 of being resampled with a probability of 0.1 of being active). The remaining eight were atypical or low-dominance exemplars in which many more features were altered (resampling probability of 0.4). The effect of this manipulation is simply to make all exemplars in a category cluster around the prototype, with high-dominance exemplars more similar to the category prototype than low-dominance exemplars. Words will be considered semantically related if they were generated from the same prototype.

The resulting 128 semantic representations were randomly assigned orthographic representations consisting of patterns of activity over 20 orthographic units. These patterns were generated randomly such that each unit had a probability of 0.1 of being active, with the constraint that every pattern had at least two active units, and all pairs of patterns differed in the activities of at least two units. No attempt was made to model orthographic relatedness among words; the orthographic patterns simply guaranteed that the written forms of words were fairly sparse and were discriminable from each other. For the current purposes, the critical property of this artificial task is that, although there are systematic relationships among word meanings, there is no systematic relationship between the written form of a word and its meaning.

Within each dominance class of each category, half of the words were designated as high-frequency and the other half as low-frequency. Each word was also assigned a single associated word, under the constraints that 1) every word was the associate of some other word; 2) associated words were never semantically related (i.e., in the same category); and 3) there were no mutual associations among word pairs (i.e., no two words were each other’s associate). The frequency of a word and its association influenced how it was selected for presentation during training, as described below.

Network Architecture. The network used to perform the task is depicted in Figure 1. The 20 orthographic units are fully connected to 100 hidden units which, in turn, are fully connected to the 100 semantic units. The semantic units themselves are fully interconnected (without self-connections) and




Figure 1: The architecture of the network. Arrows represent full connectivity between or within groups of units.

also send connections back to the hidden units. Thus, the hidden and semantic units interact in processing a given orthographic input. Including the bias terms for the hidden and semantic units, the network has a total of 32,100 connections. The weights on these connections were initialized to random values between ± 0.25 .

The states of units in the network change smoothly over time in response to influences from other units. For the purposes of simulation on a digital computer, it is convenient to approximate continuous units with finite difference equations, in which time is discretized into *ticks* of some duration τ . Thus, the input to unit j at time t is given by

$$x_j^{[t]} = \tau \sum_i s_i^{[t-\tau]} w_{ij} + (1 - \tau) x_j^{[t-\tau]} \quad (1)$$

where s_i is the state of unit i and w_{ij} is the weight from unit i to unit j . According to this equation, a unit's input at each time tick is a weighted average of its current input and that dictated by other units, where τ is the weighting proportion. A relatively large value of τ is used during most of training (0.2 in the current simulation), when minimizing computation time is critical, whereas a much smaller τ is used during testing (e.g., 0.01), when a more accurate approximation of the underlying continuous system is desired.

The state s_j of unit j at time t is simply the standard logistic or sigmoid function of its current input,

$$s_j^{[t]} = \sigma(x_j^{[t]}) = \frac{1}{1 + \exp(-x_j^{[t]})} \quad (2)$$

where $\exp(\cdot)$ is the exponential function.

Training Procedure. The network was trained in the following way. A word was presented to the network by clamping the states of the orthographic units to its assigned representation, distorted by a slight amount of random gaussian noise (with mean 0.0 and *SD* 0.05). On most trials, all other units retained the inputs and states they had at the end of processing the previous word. However, for the very first word, and with a probability of 0.01 throughout training, these units were given reinitialized inputs of 0.0 and states of 0.2. Then, for every time tick t of duration $\tau = 0.2$ over a total of 4.0 units of time, units in the network updated their states according to Equations 1 and 2. (Note that the absolute time scale of the network is arbitrary.) A continuous version of back-propagation through time (Pearlmutter, 1989) was used to calculate changes to the connection weights that would reduce the discrepancy (measured using *cross-entropy*; see Hinton, 1989) between activations of the semantic units over

the last 2.0 units of time and their correct activations for the presented word. This temporally extended error signal pressures the network to settle to the correct pattern as quickly as possible. After each word presentation, the weights were updated immediately (with a learning rate $\epsilon = 0.005$ and momentum $\alpha = 0.8$) and the next word was chosen and presented. With a probability of 0.2, the next word chosen was the associate of the previous word. On the remaining trials, the probability that words were selected for training depended on their assigned frequency, such that high-frequency words were twice as likely to be trained as low-frequency words.

After 50,000 word presentations, τ was reduced from 0.2 to 0.05, and after 3000 more presentations it was reduced to 0.01 for a final 2000 presentations. At this point, the network was completely accurate in settling into the semantic representation of each word, regardless of the preceding context.

Testing Procedure. The reaction time (RT) of the trained network in processing a word was defined as the time it took the network to settle to the point where no semantic unit changed its state by more than an *output change tolerance* of 0.001. Priming in the network occurs because this settling time is influenced by the nature of the preceding prime word(s) processed by the network. Typically, when testing the network, words are presented in prime-target pairs. First, the network is initialized to inputs of 0.0 and states of 0.2. Then the prime word is presented (with no noise) and processed for some duration (the stimulus onset asynchrony, or SOA). At this point, the target replaces the prime and the settling time of the network in response to the target is measured.

Primes can either be semantically related, associatively related, or unrelated to the target. To obtain the most reliable estimates of RT means for these various conditions, the RT for each word as target is measured when preceded by every other word as prime. For each item, its RT when preceded by each of the 15 other words in its category is averaged to yield a semantically related RT mean. Similarly, its RTs when preceded by the 111 words that are neither semantically nor associatively related yields an unrelated RT mean. Finally, the item's RT when preceded by the single word for which it is the associate is used as the associatively related RT value. These RT means were calculated at SOAs of 0.25, 0.5, 1.0, 1.5, 2.0, 2.5, 3.0, and 4.0. The manipulation of SOA is intended primarily to illustrate effects in network—direct comparisons with empirical results are problematic because longer SOAs are thought to introduce contaminating effects from subject-generated expectancies and post-lexical checking (Balota & Chumbley, 1984; Seidenberg et al., 1984).

Results and Discussion

As in other distributed network models (e.g., Masson, 1995), the distribution of RT values produced by the network in response to targets takes the form of a skewed gaussian, similar to that found in empirical studies (e.g., Ratcliff & Murdock, 1976). For instance, in processing the 128 words in a neutral context (i.e., units initialized to inputs of 0.0 and states of 0.2), the network produces RTs with mean 2.26, *SD* 0.294, and skew 0.361. The network settles faster for high-frequency words (mean 2.17) than for low-frequency words (mean 2.36; $F_{1,126}=16.35$, $p<.001$), as is typically found in word-recognition studies (e.g., Forster & Chambers, 1973).

Effect of Prime Duration on Associative Priming

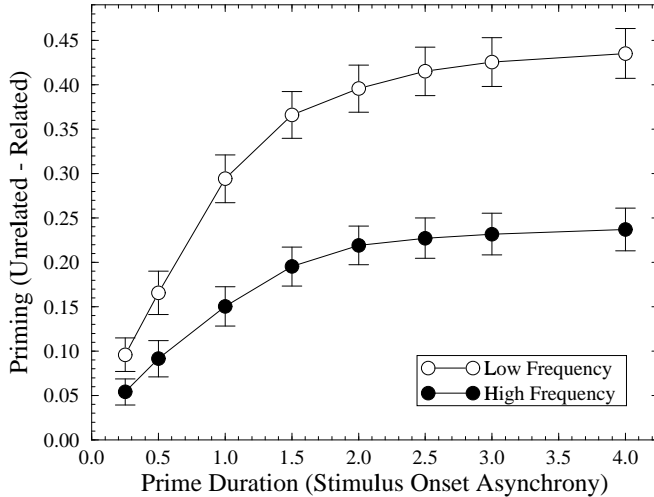


Figure 2: Effect of the duration of the prime on associative priming for low- and high-frequency targets.

However, we are more interested in differences in RTs as a function of prime-target relatedness. Thus, the remainder of the paper will report data on analyses of differences between related vs. unrelated RT means, first for associatively related words and then for semantically related words.

Associative Priming. For each word (e.g., BUTTER), the difference between the mean RTs for unrelated primes (e.g., DOG) and the RT for its associated prime (e.g., BREAD) at each SOA was computed. Positive differences reflect associative priming—faster settling for related than unrelated primes. These difference scores were entered into a $2 \times 2 \times 8$ ANOVA over items, with target category dominance (high vs. low) and frequency (high vs. low) as between-item factors and the eight values of SOA as a within-item factor. This analysis revealed a strong effect of associative priming, with a RT difference of 0.250 between unrelated versus related primes ($F_{1,124}=271.0$, $p<.001$). The degree of priming was affected strongly by SOA, with greater priming at longer SOAs ($F_{7,868}=154.8$, $p<.001$). It was also influenced by the frequency of the target, with low-frequency targets showing greater priming (mean 0.324) than high-frequency targets (mean 0.176; $F_{1,124}=23.88$, $p<.001$), as found in empirical studies (see Neely, 1991). Furthermore, frequency interacted reliably with SOA such that the difference in priming between low- and high-frequency targets increased with longer SOAs ($F_{7,868}=13.61$, $p<.001$). By contrast, there was no main effect of target category dominance ($F<1$), nor did this factor interact with frequency or SOA. Accordingly, the data were collapsed across target category dominance. Figure 2 presents the mean associative priming found across SOA for low- and high-frequency targets.

Associative priming occurs in the network because, during training, it was pressured to learn to make a rapid transition from the meaning of the prime to the meaning of its associated word much more frequently than transitions to the meanings of other words. High-frequency targets benefit less from this extra support because semantic units are already being driven more strongly than for low-frequency words; due

Effect of Prime Duration on Semantic Priming

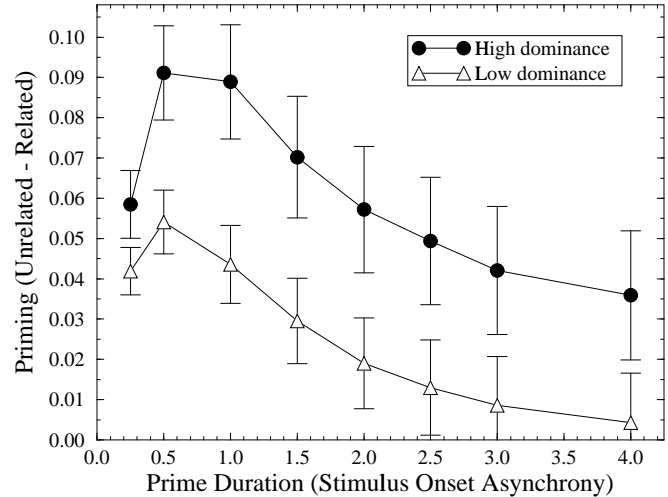


Figure 3: Effect of the duration of the prime on semantic priming for high- and low-dominance targets.

to the asymptotic behavior of the logistic activation function, any additional input to a unit yields diminishing changes in its activation as it approaches the extremes of 0.0 or 1.0.

An equivalent analysis with the associated primes and targets reversed revealed no significant backward associative priming nor any interactions ($F<1$ for all comparisons).

Semantic Priming. A similar analysis was performed on the differences in RTs for unrelated primes versus semantically related primes (e.g., BREAD–CAKE). The pattern of results is quite different from that for associative priming. Semantic priming is much weaker (mean 0.044; $F_{1,24}=29.96$, $p<.001$), and it interacts with target category dominance (means: high 0.062, low 0.027; $F_{1,124}=4.67$, $p=.033$) rather than frequency ($F<1$). Furthermore, the change in semantic priming over SOA ($F_{7,868}=22.82$, $p<.001$) is very different than for associative priming (see Figure 3 and compare with Figure 2, noting the scale difference). In particular, semantic priming peaks at very short SOAs then gradually declines as the prime is processed more fully. At intermediate SOA values, such priming may even be too weak to detect experimentally (as in the Shelton & Martin, 1992, study).

Why does semantic priming behave so differently compared with associative priming? Early on in processing a semantically related prime, units move towards a pattern that is similar to the pattern for the target, and this benefits subsequent processing of the target as long as units are still within the linear range of the logistic activation function. However, with additional processing, units are driven to more extreme values, including those that *differ* between the prime and target patterns. In order to identify the target accurately, *all* of these differences must be corrected. The time it takes for processing of the target to accomplish this is influenced by the number of units the prime and target have in common (as indicated by the effect of category dominance, matching the empirical findings), but this influence is relatively weak. As the prime is processed more fully, the time it takes to change the states of incorrect units is relatively independent of the number of other, correct units, so that semantically related

Effect of Target Degradation on Associative Priming

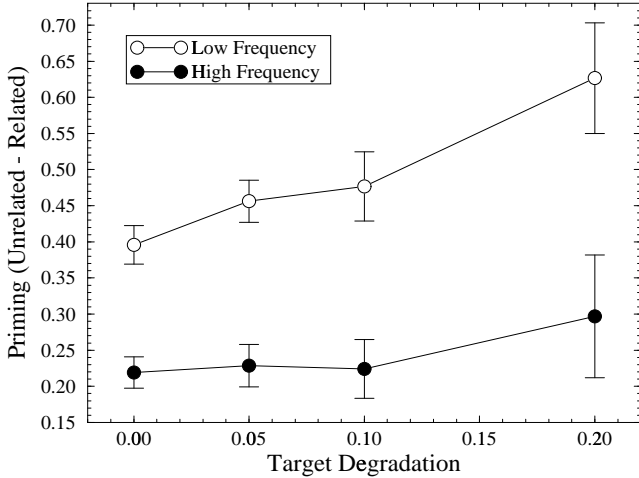


Figure 4: Effect of the degree of degradation of the target on associative priming for low- and high-frequency targets.

primes provide almost no advantage over unrelated primes.

Target Degradation. Empirical studies have found that priming is increased if targets are degraded visually (e.g., by reducing contrast; see Neely, 1991). To investigate this effect in the model, the orthographic input patterns for targets were reduced in visual contrast by scaling the input values towards the neutral value of 0.2 by varying amounts (0.05, 0.10, 0.15, and 0.2). For example, for a normal input of 1.0 and a degradation of 0.05, the presented input value would be $1.0 - 0.05(1.0 - 0.2) = 0.96$. Reaction times to degraded targets were measured when preceded by each word as a prime with an SOA of 2.0. The differences between RTs for targets preceded by their associatively related prime and their mean RTs when preceded by unrelated primes were entered into a 2x4 ANOVA over items, with target frequency as a between-item factor and target degradation as a within-item factor. Priming was influenced reliably by word frequency (means: low 0.242, high 0.489; $F_{1,126}=20.41, p<.001$), replicating the effect found with non-degraded targets. More importantly, there was greater priming for more highly degraded targets ($F_{3,378}=6.683, p<.001$; see Figure 4). The interaction of frequency and degradation was not reliable ($F_{3,378}=1.519, p=.209$). Thus, the network replicates the empirical finding that target degradation increases associative priming.

Priming Across Intervening Unrelated Items. The final experiment tested the conditions under which associative priming spanned an intervening unrelated item (e.g., BREAD-DOG-BUTTER). Such priming is observed empirically when items are processed relatively briefly. For example, Joordens and Besner’s (1992) subjects named each stimulus item under a very short (300 msec) inter-trial interval, producing naming latencies about 150 msec faster than in typical naming studies. These testing conditions can be approximated with the model by having it process the prime, the intervening item, and the target using a less-stringent criterion for the degree to which the output must settle before the network responds. Accordingly, the magnitude of associative priming across an unrelated item was investigated across increasing values of the

Associative Priming Spanning Unrelated Item

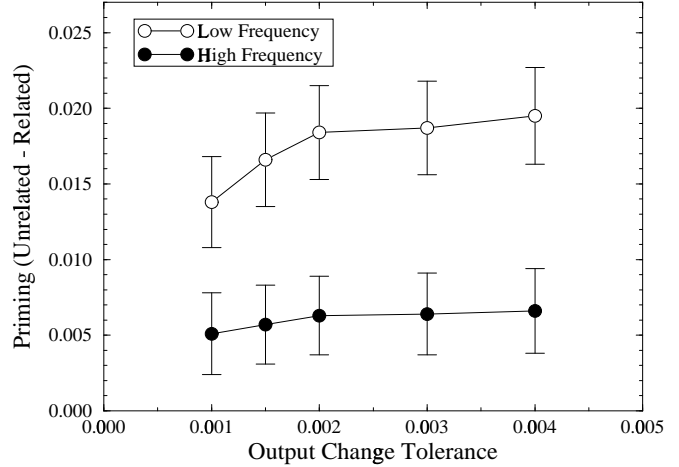


Figure 5: Effect of output change tolerance on the magnitude of associative priming across an intervening unrelated item for low- and high-frequency targets.

output change tolerance: 0.001 (the original value), 0.0015, 0.002, 0.003, and 0.004. To get a good estimate of the associatively related RT means, each association was tested with every other word in the corpus that was unrelated associatively and semantically to both the prime and target as the intermediate item. The RT means for unrelated primes were based on presenting every unrelated word as prime five times, each with a different randomly selected unrelated word as the intervening item. Responses in which the state of any semantic unit was not on the correct side of 0.5 (no more than 1.3% of trials in any condition) were considered errors and were excluded from the calculation of the means. The differences in RT means between related and unrelated prime-target pairs were then subject to a 2x5 ANOVA over items, with target frequency as a between-item factor and output change tolerance as a within-item factor. Figure 5 shows the degree of priming across an unrelated item for high- and low-frequency targets. There is a small but reliable overall priming effect (mean 0.012, $F_{1,124}=33.5, p<.001$) which is influenced both by frequency ($F_{1,124}=7.91, p=.006$) and by output change tolerance ($F_{4,496}=7.77, p<.001$), and these factors interact ($F_{4,496}=2.61, p=.035$). Thus, the network exhibits significant associative priming across an intervening unrelated item, particularly under conditions which encourage fast responding (as found by Joordens & Besner, 1992).

Conclusions

The current paper presents a distributed attractor network trained with recurrent back-propagation on an abstract version of the task of deriving the meanings of written words. In the task, semantically related words are defined to overlap in their semantic features, whereas associatively related words are defined to follow each other often during training (also see Moss et al., 1994). The network exhibits two empirical effects that have posed problems for distributed network theories: much stronger associative priming than semantic priming (Shelton & Martin, 1992), and significant associative priming across an intervening unrelated item (Joordens &

Besner, 1992). It also reproduces the empirical findings of greater priming for low-frequency targets, degraded targets, and high-dominance category exemplars (see Neely, 1991). One phenomenon it fails to reproduce, however, is *mediated priming* (e.g., LION–STRIPES, via TIGER). Thus, on the current approach, mediated priming must be attributed to weak direct associative or semantic priming (McKoon & Ratcliff, 1992). Nonetheless, the current simulation demonstrates that distributed network theories of semantic memory can account for a wide range of empirical findings on semantic and associative priming.

Acknowledgments

I thank Marlene Behrmann for helpful comments on this paper. Financial support for this research is provided by the National Institute of Mental Health (Grant MH47566) and the McDonnell-Pew Program in Cognitive Neuroscience (Grant T89–01245–016).

References

- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard Press.
- Balota, D. A., & Chumbley, J. I. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. *Journal of Experimental Psychology: Human Perception and Performance*, 10, 340–357.
- Chauvin, Y. (1988). *Symbol acquisition in humans and neural (PDP) networks*. PhD thesis, University of California, San Diego.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 407–428.
- Dosher, B. A., & Rosedale, G. (1989). Integrated retrieval cues as a mechanism for priming in retrieval from memory. *Journal of Experimental Psychology: General*, 2, 191–211.
- Fischler, I. (1977). Semantic facilitation without association in a lexical decision task. *Memory and Cognition*, 5, 335–339.
- Forster, K. I., & Chambers, S. (1973). Lexical access and naming time. *Journal of Verbal Learning and Verbal Behaviour*, 12, 627–635.
- Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91, 1–67.
- Glosser, G., & Friedman, R. B. (1991). Lexical but not semantic priming in Alzheimer's disease. *Psychology and Aging*, 6(4), 522–527.
- Hinton, G. E. (1989). Connectionist learning procedures. *Artificial Intelligence*, 40, 185–234.
- Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, 93, 411–428.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Science, USA*, 79, 2554–2558.
- Joordens, S., & Besner, D. (1992). Priming effects that span an intervening unrelated word: Implications for models of memory representation and retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(3), 483–491.
- Kawamoto, A. (1988). Distributed representations of ambiguous words and their resolution in a connectionist network. In S. L. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence*. San Mateo, CA: Morgan Kaufmann.
- Lorch, R. F. (1982). Priming and search processes in semantic memory: A test of three models of spreading activation. *Journal of Verbal Learning and Verbal Behaviour*, 21, 468–492.
- Masson, M. E. J. (1991). A distributed memory model of context effects in word identification. In D. Besner, & G. W. Humphreys (Eds.), *Basic processes in reading* (pp. 233–263). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Masson, M. E. J. (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(1), 3–23.
- McKoon, G., & Ratcliff, R. (1992). Spreading activation versus compound cue accounts of priming: Mediated priming revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(6), 1155–1172.
- McNamara, T. P. (1992a). Priming and constraints it places on theories of memory and retrieval. *Psychological Review*, 99(4), 650–662.
- McNamara, T. P. (1992b). Theories of priming I: Associative distance and lag. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(6), 1173–1190.
- McNamara, T. P. (1994). Theories of priming II: Types of primes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(3), 507–520.
- McRae, K., de Sa, V., & Seidenberg, M. S. (1993). Modeling property interactions in accessing conceptual memory. *Proceedings of the 15th Annual Conference of the Cognitive Science Society* (pp. 729–734). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Moss, H. E., Hare, M. L., Day, P., & Tyler, L. K. (1994). A distributed memory model of the associative boost in semantic priming. *Connection Science*, 6(4), 413–427.
- Moss, H. E., & Marslen-Wilson, W. D. (1993). Access to word meanings during spoken language comprehension: Effects of sentential semantic context. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19(6), 1254–1276.
- Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, 89, 609–626.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner, & G. W. Humphreys (Eds.), *Basic processes in reading* (pp. 264–336). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Pearlmutter, B. A. (1989). Learning state space trajectories in recurrent neural networks. *Neural Computation*, 1(2), 263–269.
- Postman, L., & Keppel, G. (1970). *Norms of word associations*. New York: Academic Press.
- Ratcliff, R., & McKoon, G. (1981). Does activation really spread? *Psychological Review*, 88, 454–462.
- Ratcliff, R., & McKoon, G. (1988). A retrieval theory of priming in memory. *Psychological Review*, 88, 385–408.
- Ratcliff, R., & McKoon, G. (1994). Retrieving information from memory: Spreading-activation-theories versus compound-cue-theories. *Psychological Review*, 101, 177–184.
- Ratcliff, R., & Murdock, B. B. (1976). Retrieval processes in recognition memory. *Psychological Review*, 83, 190–214.
- Seidenberg, M. S., Waters, G. S., Sanders, M., & Langer, P. (1984). Pre- and postlexical loci of contextual effects on word recognition. *Memory and Cognition*, 12, 315–328.
- Sharkey, A. J., & Sharkey, N. E. (1992). Weak contextual constraints in text and word priming. *Journal of Memory and Language*, 31(4), 543–572.
- Shelton, J. R., & Martin, R. C. (1992). How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(6), 1191–1210.
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.